arXiv:1510.00920v1 [cs.IT] 4 Oct 2015

A LEAST SQUARES APPROACH FOR STABLE PHASE RETRIEVAL FROM SHORT-TIME FOURIER TRANSFORM MAGNITUDE

Tamir Bendory and Yonina C. Eldar

Department of Electrical Engineering, Technion - Israel Institute of Technology, Haifa, Israel.

ABSTRACT

We address the problem of recovering a signal (up to global phase) from its short-time Fourier transform (STFT) magnitude measurements. This problem arises in several applications, including optical imaging and speech processing. In this paper we suggest three interrelated algorithms. The first algorithm estimates the signal efficiently from noisy measurements by solving a simple least-squares (LS) problem. In contrast to previously proposed algorithms, the LS approach has stability guarantees and does not require any prior knowledge on the sought signal. However, the recovery is guaranteed under relatively strong restrictions on the STFT window. The second approach is guaranteed to recover a non-vanishing signal efficiently from noise-free measurements, under very moderate conditions on the STFT window. Finally, the third method estimates the signal robustly from noisy measurements by solving a semi-definite program (SDP). The proposed SDP algorithm contains an inherent trade-off between its robustness and the restrictions on the STFT windows that can be used.

Index Terms— phase retrieval, short-time Fourier transform, least squares estimation, semi-definite program

1. INTRODUCTION

The problem of recovering a signal from its Fourier transform magnitude arises in many areas in engineering and science, such as optics, X-ray crystallography, speech recognition and blind channel estimation [1, 2, 3, 4, 5]. This problem is called *phase retrieval* and received considerable attention recently, partly due to its strong connections with the fields of compressed sensing and sparse recovery [6, 7, 8]. We refer the reader to a contemporary survey of the phase retrieval problem in [9].

Phase retrieval is inherently an ill-posed problem. Two main approaches have been suggested to overcome the ill-posedness. The first builds upon prior knowledge on the signal's support, such as sparsity or exact knowledge of a portion of the underlying signal [10, 11, 12, 13, 14]. The second approach makes use of additional measurements. This can be performed by structured illuminations and masks [15, 16] or by measuring the magnitude of the short-time Fourier transform (STFT). In [17], it was demonstrated that for the same number of measurements, measuring the STFT magnitude can lead to better performance than measuring an oversampled discrete Fourier transform (DFT).

In this paper we consider the recovery of a signal from its STFT magnitude. This problem arises in several applications in optics and speech processing [18, 19, 20]. In [17], the authors prove that non-vanishing signals (namely, signals whose entries are all non-zero) can be recovered by an algebraic method, under mild conditions on the STFT window. However, the algorithm is highly sensitive to noise. In [21], the authors prove that most non-vanishing one-dimensional signals can be recovered by a semi-definite program (SDP) under mild conditions on the STFT window. They also demonstrate by numerical experiments that the algorithm is robust to noise. Two additional algorithms that often work well in practice are the iterative Griffin-Lim algorithm (GLA) [19] and STFT-GESPAR (for sparse signals) [8, 17], however neither have recovery guarantees.

In this paper, we suggest three algorithms for the recovery of a signal from its STFT magnitude together with recovery guarantees. All of the proposed techniques are based on observing the DFT of the STFT magnitude. Algorithm 1 is based on a simple least-squares (LS) approach. In contrast to previous works, the LS algorithm has stability guarantees. Furthermore, it can be implemented efficiently and does not require any prior knowledge on the sought signal, whereas previous methods require the knowledge of a portion of the signal [18] or that the signal is non-vanishing [17, 21]. However, these appealing properties hold under certain restrictions on the STFT windows.

To alleviate the constraints on the window, we suggest an efficient algebraic technique, Algorithm 2, to recover non-vanishing signals. The algorithm shares similarities with the method suggested in [17], however it has slightly less restrictions on the STFT window. Both Algorithms 1 and 2 can be extended naturally to higher dimensions.

Our final approach exploits a larger set of measurements in an algorithmically more complicated way using an SDP, to robustly recover arbitrary signals. Our SDP formulation imposes different constraints than the one analyzed in [21]. Simulations show that the algorithm performs well in noisy environments. It contains the flexibility to improve the signal's estimation (by relying on more observations) at the cost of more restrictions on the STFT window.

The rest of the paper is organized as follows. Section 2 formulates the problem of recovering a signal from its STFT magnitude. Section 3 presents our algorithms, proves some of their properties and discusses connections with previous works. Section 4 shows numerical simulations, validating our theoretical findings. Section 5 concludes the work and suggests future research directions.

2. PROBLEM FORMULATION

The STFT of a signal $\mathbf{x} \in \mathbb{C}^N$ is defined as

$$\mathbf{X}[m,k] = \sum_{n=0}^{N-1} \mathbf{x}[n]\mathbf{g}[mL-n]e^{-2\pi jkn/N}, \qquad (2.1)$$

This work was funded by the European Unions Horizon 2020 research and innovation programme under grant agreement ERC-BNYQ, by the Israel Science Foundation under Grant no. 335/14, and by ICore: the Israeli Excellence Center Circle of Light.

where $k = 0, \ldots, N-1$, $m = 0, \ldots, \lceil \frac{N}{L} \rceil - 1$, $\mathbf{g}[n]$ is the STFT window and L determines the separation in time between adjacent sections. In the sequel, we assume that $\mathbf{x}$ and $\mathbf{g}$ are periodically extended over the boundaries in (2.1). The STFT can be interpreted as the Fourier transform of the signal $\mathbf{x}$, multiplied by a sliding window $\mathbf{g}$. In this work we consider $L = 1$. The analysis of $L > 1$ will be left for future work.

We assume that the information we have on the signal is the magnitude of its STFT

$$\mathbf{Y}[m,k] = |\mathbf{X}[m,k]|^2 + \eta[m,k], \tag{2.2}$$

where $\eta[m,k]$ is an additive noise. Our goal is to recover $\mathbf{x}$ given $\mathbf{Y}$. Clearly, any signal $\mathbf{z}$ and $\mathbf{z}e^{j\phi}$, for some angle ϕ, have the same STFT magnitude. Therefore, the global phase ϕ cannot be determined by any method, and we say that the recovery is *up to global phase*.

In [21], an SDP algorithm was suggested to estimate a non-vanishing signal $\mathbf{x}$ from $\mathbf{Y}$. However, the high computational complexity of the SDP problem prevents its practical use (see for instance [7]). Additionally, while the SDP algorithm of [21] seems to perform well in a noisy environment, it does not have stability guarantees. In this paper we suggest a simple and efficient algorithm to recover an arbitrary signal based on a LS formulation with stability guarantees by observing the DFT of the STFT magnitude.

Before proceeding, we introduce some notation. Let $\bar{\mathbf{z}}$ and $\mathbf{z}^*$ be the conjugate and the transpose conjugate of a vector $\mathbf{z}$, respectively. We say that $\mathbf{z}$ is non-vanishing if $\mathbf{z}[n] \neq 0$ for all n. Let $\mathbf{F}$ and $\mathbf{F}^*$ be the DFT and inverse DFT matrices. The set $\mathcal{H}_N$ denotes all Hermitian matrices of size $N \times N$, and $|\Lambda|$ the cardinality of the set Λ. For $k = 0, 1, \ldots, N-1$, we denote by $\mathrm{diag}(\mathbf{X}, k) \in \mathbb{C}^N$ the k^{th} circular diagonal of a matrix $\mathbf{X}$, namely, the entries $(i, (i+k) \mathrm{mod} N)$, $i = 0, \ldots, N-1$ of $\mathbf{X}$.

3. PHASE RETRIEVAL FROM STFT MAGNITUDE

3.1. Stable Recovery via Least-Squares Estimation

We begin by considering the DFT of $\mathbf{Y}[m,k]$ with respect to the second variable:

$$\mathbf{Z}[m,\ell] = \sum_{k=0}^{N-1} \mathbf{Y}[m,k] e^{-2\pi jk\ell/N}. \tag{3.1}$$

In the noiseless case we get

$$\begin{aligned}
\mathbf{Z}[m,\ell] &= \sum_{k=0}^{N-1} |\mathbf{X}[m,k]|^2 e^{-2\pi jk\ell/N} \\
&= \sum_{n_1,n_2=0}^{N-1} \mathbf{x}[n_1] \overline{\mathbf{x}[n_2]} \mathbf{g}[m-n_1] \overline{\mathbf{g}[m-n_2]} \sum_{k=0}^{N-1} e^{-2\pi j(n_1-n_2+\ell)k/N} \\
&= N \sum_{n=0}^{N-1} \mathbf{x}[n] \overline{\mathbf{x}[n+\ell]} \mathbf{g}[(m-n) \mathrm{mod} N] \overline{\mathbf{g}[(m-n-\ell) \mathrm{mod} N]},
\end{aligned}$$

where the last equality follows from the identity $\sum_{k=0}^{N-1} e^{-2\pi jkn/N} = N\delta[n \mathrm{mod} N]$. Consequently, for any fixed $\ell = 0, \ldots, N-1$, we obtain the linear system of equations

$$\frac{1}{N} \mathbf{z}_\ell = \mathbf{G}_\ell \mathbf{x}_\ell, \tag{3.2}$$

where $\mathbf{z}_\ell := \{\mathbf{Z}[m,\ell]\}_{m=0}^{N-1}$, $\mathbf{x}_\ell := \left\{\mathbf{x}[n] \overline{\mathbf{x}[n+\ell]}\right\}_{n=0}^{N-1}$ and $\mathbf{G}_\ell$ is an $N \times N$ circulant matrix with first column given by $\mathbf{g}[m]\overline{\mathbf{g}[(m-\ell) \mathrm{mod} N]}$.

Relation (3.2) will be the basis for the three algorithms we propose below. Let $\mathbf{X} = \mathbf{x}\mathbf{x}^*$ and observe that $\mathrm{diag}(\mathbf{X}, \ell) = \mathbf{x}_\ell$. Therefore, in the noiseless case, if the matrices $\mathbf{G}_\ell$ are invertible for all $\ell = 0, \ldots, N-1$, then we can compute $\mathbf{x}_\ell = \frac{1}{N} \mathbf{G}_\ell^{-1} \mathbf{z}_\ell$. Since $\mathbf{G}_\ell$ is a circulant matrix, it can be diagonalized by $\mathbf{G}_\ell = \mathbf{F}^* \mathbf{\Sigma}_\ell \mathbf{F}$, where $\mathbf{\Sigma}_\ell$ is a diagonal matrix whose entries are given by the DFT of the first column of $\mathbf{G}_\ell$. Thus, $\mathbf{x}_\ell = \frac{1}{N} \mathbf{F}^* {\mathbf{\Sigma}_\ell}^{-1} \mathbf{F} \mathbf{z}_\ell$. This step can be computed efficiently using FFT. The set of vectors $\mathbf{x}_\ell$, $\ell = 0, \ldots, N-1$ determines the rank one matrix $\mathbf{X}$.

In the presence of noise, we aim to solve the problem

$$\min_{\mathbf{X} \in \mathbb{C}^{N \times N}} \sum_{\ell=0}^{N-1} \|\mathbf{z}_\ell - \mathbf{G}_\ell \mathrm{diag}(\mathbf{X}, \ell)\|_2^2 \quad \text{subject to} \quad \mathrm{rank}(\mathbf{X}) = 1.$$

In our first approach, we first ignore the rank-one constraint and solve a LS problem. If the matrices $\mathbf{G}_\ell$ are invertible, then the solution is given by

$$\mathrm{diag}(\mathbf{X}, \ell) = \frac{1}{N} \mathbf{G}_\ell^{-1} \mathbf{z}_\ell = \frac{1}{N} \mathbf{F}^* {\mathbf{\Sigma}_\ell}^{-1} \mathbf{F} \mathbf{z}_\ell.$$

We then determine $\mathbf{x}$ by computing the best rank one approximation of $\mathbf{X}$, which is given by the eigenvector of $\mathbf{X}$ associated with the largest eigenvalue. This approach is summarized in Algorithm 1.

In a noise-free environment, Algorithm 1 results in exact recovery under the following conditions:

Theorem 3.1. *Any complex signal can be recovered (up to global phase) from its STFT magnitude with $L = 1$ by Algorithm 1 if the DFT of $\mathbf{g}[m]\mathbf{g}[(m-\ell) mod N]$ is non-vanishing for all $\ell = 0, \ldots, N-1$.*

Proof. In the noiseless case, Algorithm 1 recovers the $\mathbf{x}$ (up to global phase) if $\mathbf{G}_\ell$ is invertible for all $\ell = 0, \ldots, N-1$. Since $\mathbf{G}_\ell$ is a circulant matrix, it is invertible if the DFT of its first column, given by $\mathbf{g}[m]\mathbf{g}[(m-\ell) \mathrm{mod} N]$, is non-vanishing. □

In contrast to the algorithms suggested in [17, 21], Algorithm 1 is guaranteed to achieve a stable estimation in the presence of noise as the solution of a LS problem as long as the matrices $\mathbf{G}_\ell$ are invertible, without any prior knowledge on the sought signal. The algorithm is efficient due to the circular structure of the matrices $\mathbf{G}_\ell$ that can be inverted by FFT.

The main drawback of Algorithm 1 is the requirement for invertibility of $\mathbf{G}_\ell$ for all $\ell = 0, 1, \ldots, N-1$. A common choice of the STFT window is a rectangular window $\mathbf{g}[n] = \mathbf{1}_{[0,W-1]}$, with ones in the interval $[0, W-1]$ and zero elsewhere. Conditions for invertibility in this case are given by the following corollary:

Corollary 3.2. *Algorithm 1 recovers $\mathbf{x}$ (up to global phase) from its STFT magnitude with a rectangular window $\mathbf{1}_{[0,W-1]}$ and $L = 1$, if:*

1. *$N/2 < W < N$.*
2. *N and the set $\{2W-N, \ldots, W\}$ are coprime numbers. This always holds if N is a prime number.*

Proof. If $W = N$, then the DFT of $\mathbf{g}[m]\mathbf{g}[(m-\ell) \mathrm{mod} N]$ is a Delta function for all ℓ and thus vanishes. If $W \leq N/2$ then there exists ℓ for which $\mathbf{g}[m]\mathbf{g}[(m-\ell) \mathrm{mod} N] = 0$ and thus $\mathbf{G}_\ell$ is not invertible for all ℓ. If $W > N/2$ then $\mathbf{g}[m]\mathbf{g}[(m-\ell) \mathrm{mod} N]$ is a rectangular

window of length ranging between $2W - N$ and W (as a function of ℓ). Therefore, the DFT of the first column of $\mathbf{G}_\ell$ is a Dirichlet kernel (modulated by an exponential) with width ranging between $2W - N$ and W. If N and all $2W - N, \ldots, W$ are coprime then the Dirichlet kernel will not vanish. Clearly, if N is a prime number then this statement holds. □

Algorithm 1 Signal Recovery from its noisy STFT magnitude using LS

Input: The noisy magnitude of the signal STFT (2.2).
Output: Estimation of the sought signal $\hat{\mathbf{x}}$.

1. Compute $\mathbf{Z}[m, \ell]$ according to (3.1).
2. For each $\ell = 0, \ldots, N-1$, compute $\mathbf{x}_\ell = \frac{1}{N}\mathbf{F}^*\mathbf{\Sigma}_\ell^{-1}\mathbf{F}\mathbf{z}_\ell$, where $\mathbf{\Sigma}_\ell$ is a diagonal matrix with the entries $\mathbf{F}\left(\mathbf{g}[n]\mathbf{g}[(n-\ell)\mathrm{mod}N]\right)$.
3. Construct a matrix $\mathbf{X} \in \mathbb{C}^{N\times N}$ by $\mathrm{diag}(\mathbf{X}, \ell) = \mathbf{x}_\ell$ for all $\ell = 0, \ldots, N-1$.
4. Compute $\hat{\mathbf{x}} = \sqrt{\lambda_{\max}}u_{\max}$, where $\lambda_{\max}$ is the maximal eigenvalue of $\mathbf{X}$ and $u_{\max}$ is the associated eigenvector.

3.2. Algebriac Recovery in a Noise-Free Environment

To reduce the requirements on the window, suppose that the information we have is exactly $|\mathbf{X}[m, k]|^2$ without any additional noise. Theorem 3.3 states that one can recover efficiently any non-vanishing signal $\mathbf{x} \in \mathbb{C}^N$ from its STFT magnitude via Algorithm 2 if the STFT window $\mathbf{g}$ satisfies moderate conditions. Compared to Theorem 1 in [17], Theorem 3.3 does not have explicit restrictions on the size of the window, however it has an additional requirement that the DFT of $\mathbf{g}[m]\overline{\mathbf{g}[(m-\ell)\mathrm{mod}N]}$ for $\ell = 0, 1$, is non-vanishing. Theorem 4.2 in [21] has similar recovery guarantees, however it requires solving an SDP, which has high computational complexity.

Algorithm 2 Signal Recovery from its noise-free STFT magnitude

Input: The magnitude of the signal STFT $|\mathbf{X}[m, k]|^2$.
Output: Estimate of the sought signal $\hat{\mathbf{x}}$.

1. Compute $\mathbf{Z}[m, \ell]$ according to (3.1).
2. For $\ell = 0, 1$, compute $\mathbf{x}_\ell = \frac{1}{N}\mathbf{F}^*\mathbf{\Sigma}_\ell^{-1}\mathbf{F}\mathbf{z}_\ell$, where $\mathbf{\Sigma}_\ell$ is a diagonal matrix with entries $\mathbf{F}_N\left(\mathbf{g}[n]\overline{\mathbf{g}[(n-\ell)\mathrm{mod}N]}\right)$ and set $\hat{\mathbf{x}}[0] = \sqrt{\mathbf{x}_0[0]}$.
3. Compute recursively $\overline{\hat{\mathbf{x}}[n+1]} = \frac{\mathbf{x}_1[n]}{\hat{\mathbf{x}}[n]}$.

Theorem 3.3. *Any non-vanishing complex signal can be recovered (up to global phase) from $|\mathbf{X}[m, k]|^2$ with $L = 1$ by Algorithm 2 if the DFT of $\mathbf{g}[n]\mathbf{g}[(n-\ell)mod N]$ is non-vanishing for $\ell = 0, 1$.*

Proof. According to (3.2), if $\mathbf{G}_\ell$ is invertible then we get the linear system of equations

$$\frac{1}{N}\mathbf{z}_\ell = \mathbf{G}_\ell\mathbf{x}_\ell \Leftrightarrow \mathbf{x}_\ell = \frac{1}{N}\mathbf{F}^*\mathbf{\Sigma}_\ell^{-1}\mathbf{F}\mathbf{z}_\ell,$$

where Σ_ℓ is a diagonal matrix with entries $F_N\left(\mathbf{g}[n]\overline{\mathbf{g}[(n-\ell)\mathrm{mod}N]}\right)$. By assumption, the DFT of $|\mathbf{g}[n]|^2$ is not equal to zero and thus we can compute $\mathbf{x}_0[n] = |\mathbf{x}[n]|^2$. If the signal is known to be non-negative, then the algorithm is completed. By further assumption, we can compute $\mathbf{x}_1[n] = \mathbf{x}[n]\overline{\mathbf{x}[n+1]}$. Since we cannot recover the global phase, we arbitrary set $\mathbf{x}[0] = |\mathbf{x}[0]|$ which is known and equal to $\sqrt{\mathbf{x}_0[0]}$. Now, we can compute the rest of the entries recursively according to step 3. □

Algorithm 2 alleviates the requirements of the Theorem 3.1 on the window. The following corollary demonstrates this on a rectangular window. The proof is similar to the proof of Corollary 3.2 and thus omitted.

Corollary 3.4. *Algorithm 2 recovers a non-vanishing $\mathbf{x}$ (up to global phase) with a rectangular window $\mathbf{g}[n] = \mathbf{1}_{[0,W-1]}$ and $L = 1$, if:*

1. $2 \le W \le N - 1$.
2. *N and $W - \ell$ are coprime for $\ell = 0, 1$.*

3.3. Extension to High-Dimensional Signals

We briefly describe the extension of Algorithms 1 and 2 to bivariate signals. The algorithms can be extended to higher dimensions by the same methodology. The two-dimensional (2D) STFT with respect to a bivariate window $\mathbf{g}$ is given by (for $L = 1$)

$$\mathbf{X}[m_1, m_2; k_1, k_2] = \sum_{n_1,n_2=0}^{N-1} \mathbf{x}[n_1, n_2]\,\mathbf{g}[m_1 - n_1, m_2 - n_2]\, e^{-2\pi j(k_1 n_1 + k_2 n_2)/N}.$$

In the same manner as in (3.1), we take the 2D DFT of $|\mathbf{X}[m_1, m_2; k_1, k_2]|^2$ with respect to k_1 and k_2 and get, [1]

$$\begin{aligned}
&\frac{1}{N^2}\mathbf{Z}[m_1, m_2; \ell_1, \ell_2] \\
&= \frac{1}{N^2}\sum_{k_1,k_2=0}^{N-1} |\mathbf{X}[m_1, m_2; k_1, k_2]|^2 e^{-2\pi j(k_1\ell_1 + k_2\ell_2)/N} \\
&= \sum_{n_1,n_2=0}^{N-1} \mathbf{x}[n_1, n_2]\,\overline{\mathbf{x}[n_1 + \ell_1, n_2 + \ell_2]} \\
&\quad \mathbf{g}[m_1 - n_1, m_2 - n_2]\,\overline{\mathbf{g}[m_1 - n_1 - \ell_1, m_2 - n_2 - \ell_2]}.
\end{aligned}$$

For a fixed (ℓ_1, ℓ_2), we obtain the linear system of equations

$$\frac{1}{N^2}\mathbf{Z}[m_1, m_2; \ell_1, \ell_2] = \mathbf{G}_{\ell_1,\ell_2}\mathbf{x}_{\boldsymbol{\ell_1},\boldsymbol{\ell_2}}, \tag{3.3}$$

where $\mathbf{x}_{\boldsymbol{\ell_1},\boldsymbol{\ell_2}} = \mathbf{x}[n_1, n_2]\,\overline{\mathbf{x}[n_1 + \ell_1, n_2 + \ell_2]}$ and $\mathbf{G}_{\ell_1,\ell_2}$ is a block circulant matrix with circulant blocks. Such a matrix is diagonalized by the 2D DFT matrix and the eigenvalues are given by the 2D DFT of the first column of $\mathbf{G}_{\ell_1,\ell_2}$. Therefore, if the 2D DFT of the first column of $\mathbf{G}_{\ell_1,\ell_2}$ is non-vanishing then the system is invertible and one can recover $\mathbf{x}_{\boldsymbol{\ell_1},\boldsymbol{\ell_2}}$ efficiently. The rest of the procedure is a simple generalization of the one-dimensional cases presented in Algorithms 1 and 2.

[1] All the indices in this section should be interpreted as modolu N.

3.4. Stable Recovery via Semidefinite Programming

Algorithm 2 uses merely two columns of the matrix $\mathbf{Y}[m,\ell]$ (3.1). Algorithm 3 suggests to exploit a larger set of columns of $\mathbf{Y}[m,\ell]$, to achieve robust estimation in the presence of noise.

Let $\Lambda \subseteq \{0,1,\ldots,N-1\}$ and let $|\Lambda|$ be its cardinality. Algorithm 3 suggests to estimate the signal from noisy observations by minimizing the trace of a matrix among all positive semidefinite Hermitian matrices which satisfy the $|\Lambda|$ constraints on the matrix diagonals. It is well-known that minimizing the trace of an Hermitian matrix promotes a low-rank solution.

We expect that choosing a larger set of observations $|\Lambda|$ will lead to a better estimation of $\mathbf{x}$. This is true if the matrices $\mathbf{G}_\ell$ are invertible for all $\ell \in \Lambda$. Therefore, Algorithm 3 enables the flexibility to choose the number of observation we use depending on the STFT window. For instance, if we consider the rectangular window $\mathbf{1}_{[0,W-1]}$ for some $W < N/2$ and the set $\Lambda = \{0,1,\ldots,\lambda-1\}$, then the matrices $\mathbf{G}_\ell$ are invertible only if $\lambda \leq W$ and that N and $W-\ell$ are coprime for all $\ell \in \Lambda$. If this is the situation, we expect that increasing λ (up to W) will result in a better estimation. The numerical results in Section 4 corroborate this conclusion.

Algorithm 3 Recovering a signal from its noisy STFT magnitude using SDP

Input: The noisy magnitude of the signal STFT (2.2).
Output: Estimate of the sought signal $\hat{\mathbf{x}}$.

1. Compute $\mathbf{Z}[m,\ell]$ according to (3.1).
2. Choose a set $\Lambda \subseteq \{0,\ldots,N-1\}$ and solve :

$$\begin{aligned} \mathbf{X} = \min_{\widetilde{\mathbf{X}} \in \mathcal{H}_N} \operatorname{trace}(\widetilde{\mathbf{X}}) \quad \text{subject to} \quad \widetilde{\mathbf{X}} \geq 0, \\ \|\mathbf{z}_\ell - \mathbf{G}_\ell \operatorname{diag}(\widetilde{\mathbf{X}},\ell)\|_2 \leq \eta_\ell, \quad \forall \ell \in \Lambda, \end{aligned} \tag{3.4}$$

   where η_ℓ is the noise level associated with the ℓ^{th} column of $\mathbf{Z}[m,\ell]$.
3. $\hat{\mathbf{x}} = \sqrt{\lambda_{\max}} u_{\max}$, where $\lambda_{\max}$ is the maximal eigenvalue of $\mathbf{X}$ and $u_{\max}$ is the associated eigenvector.

4. NUMERICAL EXPERIMENTS

This section is devoted to numerical experiments, examining the three algorithms presented above. In all experiments, both the real and the imaginary parts of the signals were drawn from an iid normal distribution. An additive noise with normal distribution was added with the appropriate variance, according to the desired signal to noise (SNR) ratio. The error was measured by $\frac{\|\mathbf{x}-\hat{\mathbf{x}}\|_2}{\|\mathbf{x}\|_2}$, where $\hat{\mathbf{x}}$ is the algorithm output.

We compared Algorithm 1 with the GLA [19] and the noisy version of the SDP algorithm in [21]. We solved the SDP using CVX [22] and stopped the iterative GLA if the difference between consecutive iterations was less that 10^{-6} or after 500 iterations. Figure 1 demonstrates that all three algorithms have similar performance for signals of length $N = 23$ and Gaussian window e^{-n^2/σ^2} with $\sigma = \lceil \frac{N}{2} \rceil$. However, the average running time of the GLA and the SDP of [21] was 27 and 2634 times greater than the running time of Algorithm 1, respectively. Table 1 presents the recovery error of Algorithm 2 in a noise-free environment. We used a rectangular window with different lengths. As predicted by Theorem 3.3, the recovery error is negligible. Figure 2 compares the recovery error of Algorithm 3 with the SDP algorithm of [21] for signals of length $N = 23$ and a rectangular window with $W = 5$ as a function of the SNR. The problems were solved by CVX. We chose $\Lambda = \{0,1,\ldots,\lambda-1\}$ for $\lambda = 3, 5, 23$. As can be seen, Algorithm 3 is stable even for small values of λ. Additionally, increasing λ beyond W does not reduce the recovery error significantly as predicted in Section 3.4. The running time of the SDP algorithm of [21] was 3.5 and 12 times greater than the running time of Algorithm 3 with $\lambda = 5, 23$, respectively.

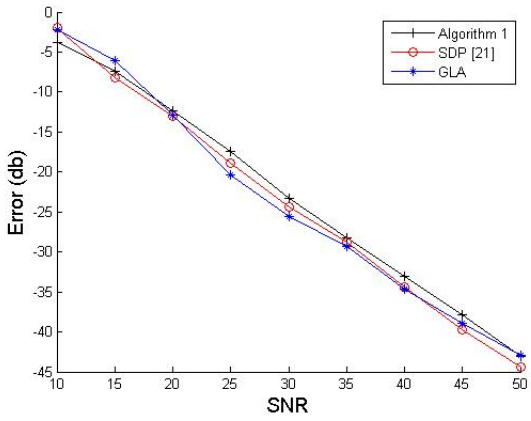

Fig. 1: The recovery error of Algorithm 1, the SDP method suggested in [21] and the GLA [19] as a function of the SNR. Each experiment was conducted 10 times.

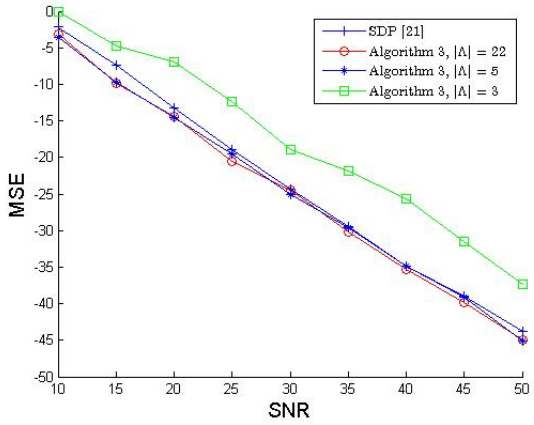

Fig. 2: The recovery error of Algorithm 3 with a rectangular window of length $W = 9$ and $N = 23$ as a function of the SNR for different values of $|\Lambda|$. The experiment was conducted 10 times for each value.

5. CONCLUSION

In this work, we suggested three interrelated algorithms to recover a signal from its STFT magnitude. Each algorithm suits a different case. Algorithm 1 is efficient and robust to noise but restricts the potential STFT windows. Algorithm 2 is an efficient and exact algorithm in the noise-free environment. Algorithm 3 uses an

	$W = 5$	$W = 23$	$W = 41$
Mean error	3.52×10^{-12}	6.84×10^{-12}	1.13×10^{-11}
Max error	1.46×10^{-12}	3.05×10^{-12}	7.02×10^{-11}

Table 1: The recovery error of Algorithm 2 with a rectangular window of length W and $N = 211$ in a noise-free environment. The experiment was conducted 100 times for each window size.

SDP formulation and contains the flexibility to control its stability to noise at the cost of restrictions on the STFT windows. Proving its robustness and properties is the goal of future work.

All the results in this paper were focused on maximal overlapping in the STFT windows, i.e. $L = 1$. Previous papers have indicated that this restriction can be alleviated. The analysis of the case $L > 1$ requires incorporating new tools into the suggested framework and is under ongoing research.

6. REFERENCES

[1] R.W Harrison, "Phase problem in crystallography," *JOSA A*, vol. 10, no. 5, pp. 1046–1055, 1993.

[2] A Walther, "The question of phase retrieval in optics," *Journal of Modern Optics*, vol. 10, no. 1, pp. 41–49, 1963.

[3] R.P Millane, "Phase retrieval in crystallography and optics," *JOSA A*, vol. 7, no. 3, pp. 394–411, 1990.

[4] B Juang and L Rabiner, "Fundamentals of speech recognition," *Signal Processing Series. Prentice Hall, Englewood Cliffs, NJ*, 1993.

[5] B Baykal, "Blind channel estimation via combining autocorrelation and blind phase estimation," *Circuits and Systems I: Regular Papers, IEEE Transactions on*, vol. 51, no. 6, pp. 1125–1131, 2004.

[6] E.J Candes, Y.C Eldar, T Strohmer, and V Voroninski, "Phase retrieval via matrix completion," *SIAM Review*, vol. 57, no. 2, pp. 225–251, 2015.

[7] E.J Candes, X Li, and M Soltanolkotabi, "Phase retrieval via wirtinger flow: Theory and algorithms," *Information Theory, IEEE Transactions on*, vol. 61, no. 4, pp. 1985–2007, 2015.

[8] Y Shechtman, A Beck, and Y.C Eldar, "Gespar: Efficient phase retrieval of sparse signals," *Signal Processing, IEEE Transactions on*, vol. 62, no. 4, pp. 928–938, 2014.

[9] Y Shechtman, Y.C Eldar, O Cohen, H.N Chapman, Jianwei J Miao, and M Segev, "Phase retrieval with application to optical imaging: a contemporary overview," *Signal Processing Magazine, IEEE*, vol. 32, no. 3, pp. 87–109, 2015.

[10] R.W Gerchberg, "A practical algorithm for the determination of phase from image and diffraction plane pictures," *Optik*, vol. 35, pp. 237, 1972.

[11] J.R Fienup, "Phase retrieval algorithms: a comparison," *Applied optics*, vol. 21, no. 15, pp. 2758–2769, 1982.

[12] Y Shechtman, Y.C Eldar, A Szameit, and M Segev, "Sparsity based sub-wavelength imaging with partially incoherent light via quadratic compressed sensing," *Optics express*, vol. 19, no. 16, pp. 14807–14822, 2011.

[13] J Ranieri, A Chebira, Y.M Lu, and M Vetterli, "Phase retrieval for sparse signals: Uniqueness conditions," *arXiv preprint arXiv:1308.3058*, 2013.

[14] K Jaganathan, S Oymak, and B Hassibi, "Sparse phase retrieval: Uniqueness guarantees and recovery algorithms," *arXiv preprint arXiv:1311.2745*, 2013.

[15] E.J Candes, X Li, and M Soltanolkotabi, "Phase retrieval from coded diffraction patterns," *Applied and Computational Harmonic Analysis*, 2014.

[16] D Gross, F Krahmer, and R Kueng, "Improved recovery guarantees for phase retrieval from coded diffraction patterns," *arXiv preprint arXiv:1402.6286*, 2014.

[17] Y.C Eldar, P Sidorenko, D.G Mixon, S Barel, and O Cohen, "Sparse phase retrieval from short-time fourier measurements," *Signal Processing Letters, IEEE*, vol. 22, no. 5, pp. 638–642, 2015.

[18] S.H Nawab, T.F Quatieri, and J.S Lim, "Signal reconstruction from short-time fourier transform magnitude," *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 31, no. 4, pp. 986–998, 1983.

[19] D.W Griffin and J.S Lim, "Signal estimation from modified short-time fourier transform," *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 32, no. 2, pp. 236–243, 1984.

[20] R Trebino, "Frequency resolved optical gating: The measurement of ultrashort optical pulses," *Kluwer Academic Publishers, Boston*, vol. 5, pp. 70–73, 2002.

[21] K Jaganathan, Y.C Eldar, and B Hassibi, "STFT phase retrieval: Uniqueness guarantees and recovery algorithms," *arXiv preprint arXiv:1508.02820*, 2015.

[22] M Grant, S Boyd, and Y Ye, "Cvx: Matlab software for disciplined convex programming," 2008.